March 2004

www.veritest.com • info@veritest.com

AMD: WebBench Virtualization Performance Study

Test report prepared under contract from Advanced Micro Devices, Inc.

Executive summary

Advanced Micro Devices, Inc. ("AMD") commissioned VeriTest, a division of Lionbridge Technologies, Inc., to conduct a competitive Web server performance study within a virtualized server environment. We conducted this study by running performance tests on a virtualized AMD Opteron™ processor-based server environment versus a virtualized Intel® Xeon™ processor-based server environment. We conducted the study by running a set of performance tests using PC Magazine WebBench 5.0.

Key findings

- ❑ At the 1 virtual machine (VM) per CPU configuration, we found that the AMD Opteron server had a peak of 2767.79 requests per second in our test environment. In the same configuration, the Intel Xeon server had a peak of 2044.87 requests per second, or 26% less than the AMD Opteron 1 VM per CPU configuration. [See below]

To perform this study, AMD supplied VeriTest with one HP ProLiant DL585 AMD Opteron processor-based server and one IBM xSeries 445 Intel Xeon processor-based server. These servers were chosen because they were considered to be near-equivalent platforms from a configuration perspective. Both servers had quad-processor configurations, 16 GB of RAM, and four Intel 82543 Gigabit Ethernet Adapters. See Appendix A for additional server specifications details.

We used VMware ESX Server 2.5.0 as the server virtualization software. Virtualization software allows one to partition a single hardware server into multiple virtual machines, each with isolated set of virtual hardware including CPU, memory, disk and networking, and running an independent operating system. Virtualization software maps these multiple independent sets of virtual hardware to underlying server hardware, such as CPU, memory, and network resources, so that each virtual machine gets a share of the underlying hardware resources.

For each server, we created three virtualized environments using ESX Server. Since both servers had four processors, we configured the servers with virtual machines (VMs) in multiples of four. The configurations that we tested were:

- 1 VM per CPU, or 4 VMs total, each VM with 3400 MB of RAM
- 2 VMs per CPU, or 8 VMs total, each VM with 1700 MB of RAM
- 3 VMs per CPU, or 12 VMs, each VM with 1100 MB of RAM

Within each VM, we installed Microsoft Windows Server 2003 Enterprise Edition.

We configured IIS 6.0 within each VM and configured a WebBench client testbed to generate the WebBench eCommerce CGI workload to all of the VMs simultaneously. The WebBench eCommerce CGI workload uses a simple CGI executable to provide dynamic content as well as test the secure transaction features of the Web server. We designed the test to distribute the WebBench workload across all of the VMs in an evenly distributed manner. We ran the WebBench eCommerce CGI workload using three engines per client to measure the requests per second, throughput, and latency of the servers in these virtualized environments.

As shown in Figure 1, below, we found that the AMD Opteron server running in a 1 VM per CPU configuration had the greatest number of requests per second, with a peak of 2767.79 requests per second in our test

environment. The Intel Xeon server running in a 1 VM per CPU configuration had a peak of 2044.87 requests per second, or 26% less than the AMD Opteron 1 VM per CPU configuration

The AMD Opteron server running in a 2 VM per CPU configuration had a peak of 2580.67 requests per second. The Intel Xeon server running in a 2 VM per CPU configuration had a peak of 1846.85 requests per second, or 28% less than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a peak of 2401.41 requests per second. The Intel Xeon server running in a 3 VM per CPU configuration had a peak of 1854.11 requests per second, or 23% less than the AMD Opteron 3 VM per CPU configuration

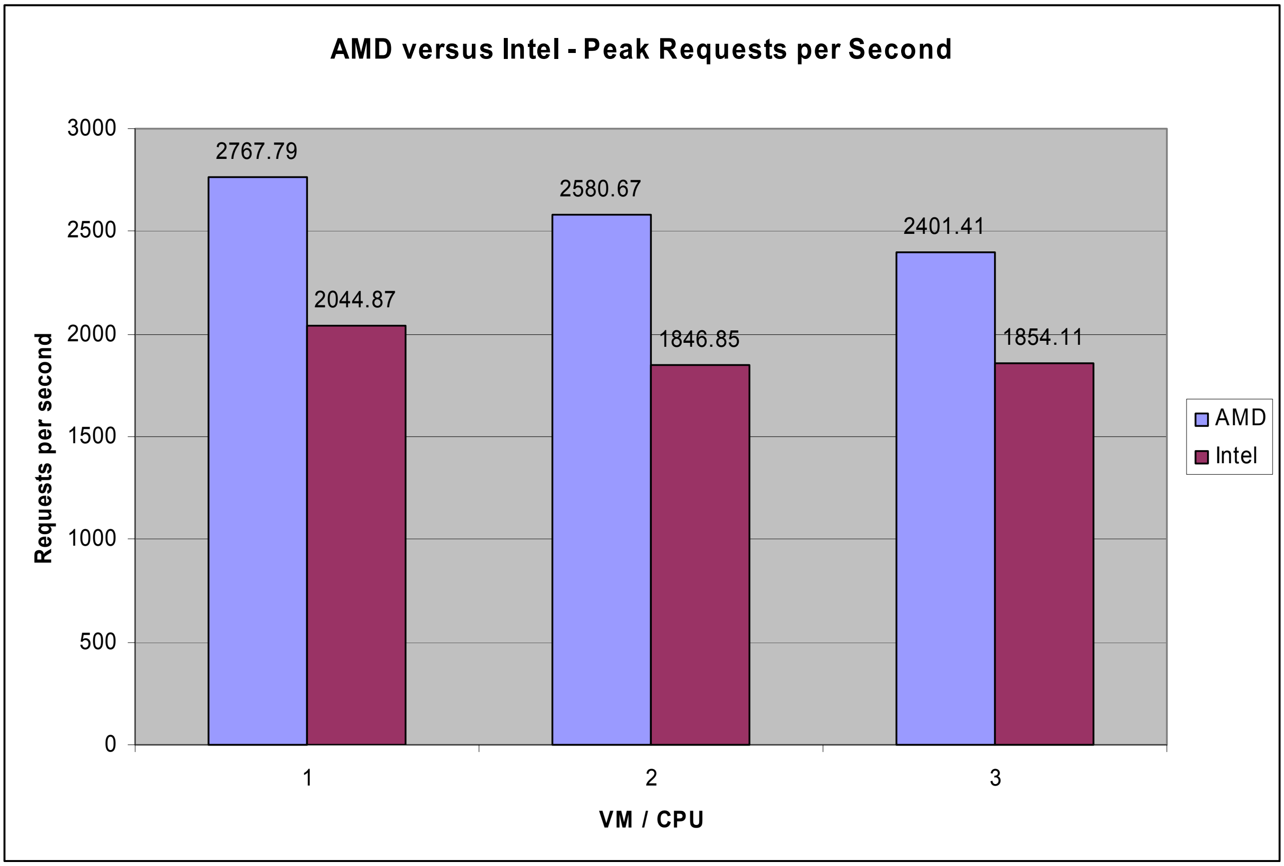

Figure 1: WebBench requests per second for AMD versus Intel in 1, 2, and 3 VM per CPU configurations (Greater numbers are better.)

We found that the AMD Opteron server running in a 1 VM per CPU configuration had the greatest WebBench peak throughput of 10.72 MB per second in our test environment. The Intel Xeon server running in a 1 VM per CPU configuration had a peak throughput of 7.94 MB per second, or 26% less than the AMD Opteron 1 VM per CPU configuration

The AMD Opteron server running in a 2 VM per CPU configuration had a peak throughput of 10.0 MB per second. The Intel Xeon server running in a 2 VM per CPU configuration had a peak throughput of 7.15 MB per second, or 28% less than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a peak throughput of 9.23 MB per second. The Intel Xeon server running in a 3 VM per CPU configuration had a peak throughput of 7.18 MB per second, or 22% less than the AMD Opteron 3 VM per CPU configuration.

Overall, we found that the AMD Opteron server had greater requests per second and throughput peaks for all of the VMs than any of the requests per second and throughput peaks on the Intel Xeon server VMs. We found that the AMD Opteron server in the 3 VM per CPU configuration had greater peak requests per second and peak throughput than the Intel Xeon server in the 1 VM per CPU configuration.

We found that the AMD Opteron server running in a 1 VM per CPU configuration had the lowest WebBench request latency of 65.6 milliseconds at the 60-client load in our test environment. The Intel Xeon server running in a 1 VM per CPU configuration had a request latency of 88.6 milliseconds at the 60-client load, or 35% higher than the AMD Opteron 1 VM per CPU configuration. Lower results are better.

The AMD Opteron server running in a 2 VM per CPU configuration had a request latency of 74.5 millisecond at the 64-client load. The Intel Xeon server running in a 2 VM per CPU configuration had a request latency of 110.4 milliseconds at the 64-client load, or 48% higher than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a request latency of 151.2 milliseconds at the 120-client load. The Intel Xeon server running in a 3 VM per CPU configuration had a request latency of 194.5 milliseconds at the 120-client load, or 29% higher than the AMD Opteron 3 VM per CPU configuration.

We found that the AMD Opteron server running in 1, 2, and 3 VM per CPU configurations had lower mean WebBench request latencies than did the Intel Xeon server running in 1, 2, and 3 VM per CPU configurations.

Testing methodology

AMD commissioned VeriTest to conduct a competitive Web server performance study within virtualized server environments. We conducted this competitive study by running performance tests on a virtualized AMD Opteron processor-based server environment versus a virtualized Intel Xeon processor-based server environment. We conducted the study by running a set of performance tests using PC Magazine's WebBench 5.0.

To perform this study, AMD supplied VeriTest with one HP ProLiant DL585 AMD Opteron processor-based server and one IBM xSeries 445 Intel Xeon processor-based server. See Appendix A for additional server specifications details.

The two servers that we used for this study were:

- One HP ProLiant DL585 server equipped with four AMD Opteron 2.40 GHz processors, 16 GB of RAM (node memory mode), 256 KB of L2 cache, 4 MB of L3 cache, four Intel 82543 Gigabit Ethernet Adapters, an HP Smart Array 5i disk controller, an HP Smart Array 6400 U320 disk controller, and two HP 73.9GB 15K Ultra320 SCSI disk drives configured as RAID 0.

- One IBM xSeries 445 server equipped with four Intel Xeon 3.00 GHz processors with Hyper-Threading enabled, 16 GB of RAM, 1 MB of L2 cache, 4 MB of L3 cache, four Intel 82543 Gigabit Ethernet Adapters, a ServerRAID-6M disk controller, and two 73.4 GB 15K Ultra320 SCSI disk drives configured as RAID 0.

Virtualization

AMD requested that we measure the performance of a Web server within a set of VMs. Since both servers in our testbed contained quad-processor configurations, to ensure a symmetric CPU affinity configuration in all test cases, we chose to scale the number of VMs by multiples of four. Based on this rule, we tested in the following configurations:

- 1 VM per CPU, or 4 virtual machines
- 2 VM per CPU, or 8 virtual machines
- 3 VM per CPU, or 12 virtual machines

To create the VMs, we used VMware ESX Server 2.5.0 build-11548 with no special ESX server tuning parameters.

We installed the VMware software on both servers by using the default installation method for ESX server. We chose to allocate 272MB RAM to the virtualization software, which supported a maximum allowable of 16 virtual machines. We defined our disk partitioning setup as shown in Figure 2.

Partition Type	Size (MB)
boot	52
/	2498
swap	1498
extended	136967
-> vmfs	136863
-> vmcore	100
-> free space	4

Figure 2: ESX Server Storage partitioning

Both servers in our testbed contained 16GB of RAM. We varied the amount of RAM available to each VM based on the total number of VMs in each test case.

- 4 virtual machines ⇨ 3400 MB per VM. Total = 13.6 GB
- 8 virtual machines ⇨ 1700 MB per VM. Total = 13.6 GB
- 12 virtual machines ⇨ 1100 MB per VM. Total = 13.2 GB

We used a virtual disk of 5GB size for each of the virtual machines in all test cases.

The VMware ESX Server software contained a feature called “Virtual Ethernet Switch” that provides the ability to assign specific VMs to specific network adapters. See Figure 3 for a diagram of how the Virtual Ethernet Switch features works.

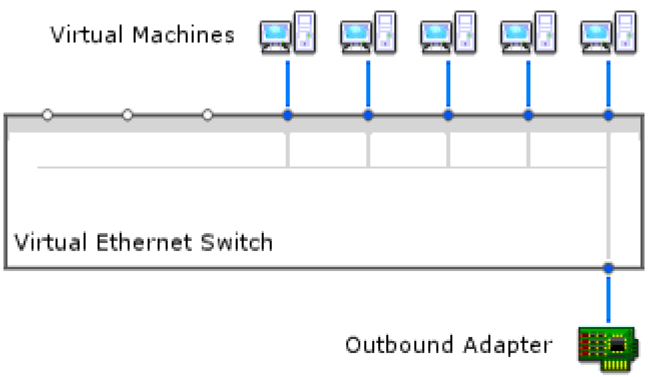

Figure 3: Virtual Ethernet Switch Configuration

To ensure we distributed the WebBench load evenly across VMs, we created a set of Port Groups to specific network adapters. We assigned Port Group numbers to specific VMs to ensure an even distribution of WebBench workload activity. We configured the Port Groups on the servers under test as listed below:

VM 1, VM5, VM9	⇨ Port Group 1	⇨ Intel 82543 Gigabit Ethernet Adapters #1
VM 2, VM6, VM10	⇨ Port Group 2	⇨ Intel 82543 Gigabit Ethernet Adapters #2
VM 3, VM7, VM11	⇨ Port Group 3	⇨ Intel 82543 Gigabit Ethernet Adapters #3
VM 4, VM8, VM12	⇨ Port Group 4	⇨ Intel 82543 Gigabit Ethernet Adapters #4

We configured WebBench to increment the client workload in an evenly distributed, VM round robin fashion. See Appendix C for a mapping of WebBench clients to VMs for the different server configurations.

As part of the VMware ESX Server installation, we installed the VMware VirtualCenter application on an additional system in our network testbed to manage the ESX Server configurations on both servers under test. We used VirtualCenter to create the original VM on both servers and then to clone the original VM multiple times. See Figure 4 for screenshot of the VMware VirtualCenter user interface.

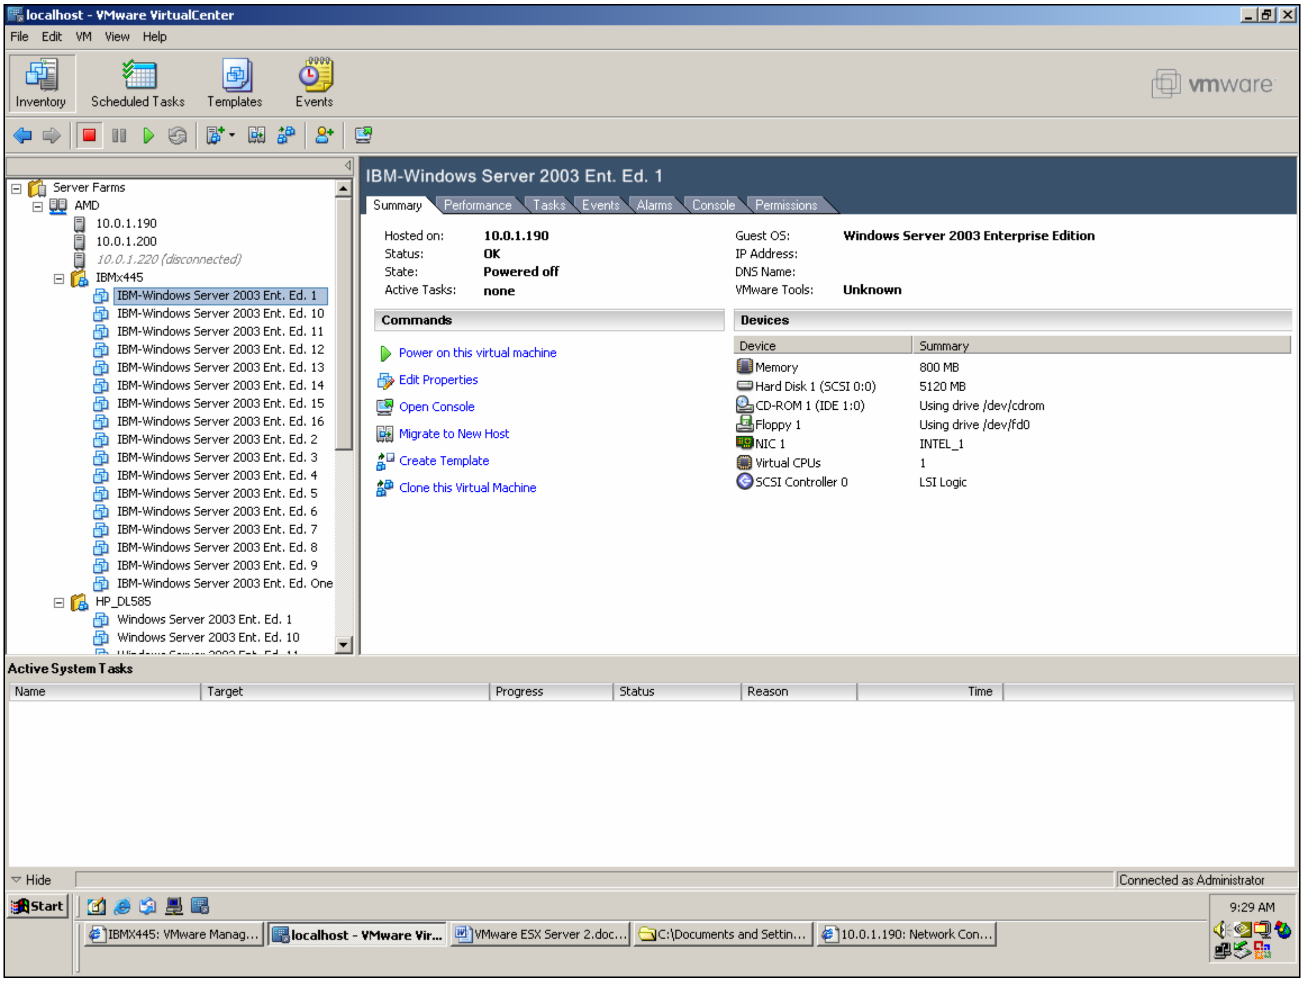

Figure 4: VMware VirtualCenter User Interface

Within each VM, we installed Microsoft Windows 2003 Enterprise Edition, version 5.2.3790 Build 3790 with no added updates or additions.

We installed the WebBench test tree within the IIS C:\Inetpub\wwwroot folder in each VM. We additionally conducted a set of IIS configuration changes as described below:

- On the “WebSite” tab, we disabled logging.
- On the “Home Directory” tab, we set “Execute Permissions” to “Scripts and Executables” for the document root directory.
- On the “Home Directory” tab, we set “Script source access”, “Read”, “Write”, and “Directory Browsing” for the document root directory.
- We added “MIME Type” of .ex/WebBench.

Additionally, we made the following registry modifications on the Web server:

- Set HKLM\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters\MaxCachedFileSize to 725000 Kilobytes.
- Set HKLM\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters\ObjectCacheTTL to 7200 seconds.

Finally, we added a server certificate to each IIS Web server.

We created one VM on each server under test as described above. We then used the Clone feature within VMware to make multiple copies of the original VM. This ensured that all VMs were identical to the original VM and ensured that we minimized performance variations due to inconsistent VM configurations.

For the IBM xSeries 445 server, we enabled HyperThreading in both the ESX Server and in hardware.

WebBench

To measure the Web server performance, we used PC Magazine's WebBench 5.0 (128-bit US version), an industry standard web server performance benchmark that measures the performance of web server software and hardware. This was done by using multiple clients making HTTP 1.0 GET requests to the web server to simulate and generate web traffic. WebBench simulates the web traffic by running the clients in various increments; these increments are referred to as a "mix". A full WebBench run consists of 16 mixes and starts with 1 client and adds 4 additional clients per mix. We used the eCommerce test suite which includes a mixture of requests for static HTML objects, dynamic CGI-based objects, and objects requiring SSL sessions. VeriTest modified the default eCommerce CGI test suite by changing the number of client engines running on each WebBench client from the default value of 1 to 3. This modification allows the WebBench clients to significantly increase the load generated by each individual test client and put a heavier load on the Web server with the same number of clients when compared to the standard WebBench test suites that utilize a single engine on each physical client.

To generate the WebBench workload for the testing, we used a total of 120 Intel processor-based Dell PowerEdge 350 servers, each configured with a single 850 MHz Pentium-III processor, 256MB RAM, a single 10GB 7,200RPM IDE hard disk drive, and dual Intel Pro/100+ Ethernet adapters running a mix of Windows 2000 and Windows XP Professional SP1 which were used as the WebBench clients. We configured the clients into four subnets of 30 WebBench clients each and then connected each subnet to its own an Extreme Networks Summit48 Fast Ethernet switch. We connected a Gigabit Ethernet port from each of the four Extreme Networks Summit48 Fast Ethernet switches to one of the four Intel 82543 Gigabit Ethernet adapters installed in the server. The Extreme switch was set to auto-negotiate line speed and duplex, resulting in full duplex Gigabit Ethernet connections between the server and the switch and full duplex Fast Ethernet connections between the clients and the switch. Refer to Appendix B for details on the test bed configuration.

For this test configuration, we had to ensure that the IIS server on each individual VM received WebBench workload in an evenly distributed manner. Since each VM behaved as a unique server, we assigned each VM a unique static IP address. We configured the WebBench testbed to add new workload requests evenly across the web servers in each VM as the WebBench load increased throughout the test.

Test cases

After we completed the system and VM configuration as described above, we ran the following test cases:

Processor type	VMs / CPU	Total VMs	RAM allocation	Number of WebBench clients
AMD Opteron	1	4	3400MB	60
AMD Opteron	2	8	1700MB	64
AMD Opteron	3	12	1100MB	120
Intel Xeon	1	4	3400MB	60
Intel Xeon	2	8	1700MB	64
Intel Xeon	3	12	1100MB	120

Figure 5: Virtualization test case configurations

As shown in Figure 5, we used 60 WebBench clients to test the 1 VM per CPU and 64 WebBench clients to test the 2 VMs per CPU because this allowed us to complete the test mix with a load that was evenly distributed across all VMs. We used 120 WebBench clients to test the 3 VMs per CPU because we needed additional Web server load to find the peak throughput and requests per second of that configuration.

While we conducted WebBench testing, we also collected server subsystem performance metrics by using the *esxtop* tool provided with the VMware software. *esxtop* collected the physical CPU utilization, the CPU utilization of virtual machines, NIC bandwidth, and disk utilization. We paid special attention to the CPU utilization during our testing. We ensured that we reached peak Web server performance by validating that physical CPU utilization of the underlying server had reached 100% when we reached the peak throughput values.

Since we tested the scaling of Web server performance, as we increased the number of active VMs in each test case, we effectively reduced the amount of Web server load associated to each specific VM. Because of this, we increased the overall Web server load as we scaled from 1 VM per CPU to 2 VMs per CPU to 3 VMs per CPU.

To illustrate the effect of the reduced workload per VM, see Appendix C. As can be seen in the tables in this appendix, at a given load level, each VM in the 1 VM per CPU test case receives more WebBench clients than do the VMs in the 3 VM per CPU test case.

Test results

AMD commissioned VeriTest to conduct a competitive Web server performance study within virtualized server environments. We conducted this competitive study by running performance tests on a virtualized AMD Opteron processor-based server environment versus a virtualized Intel Xeon processor-based server environment. We conducted the study by running a set of performance tests using PC Magazine's WebBench 5.0.

To perform this study, AMD supplied VeriTest with one HP ProLiant DL585 AMD Opteron processor-based server and one IBM xSeries 445 Intel Xeon processor-based server. Each server contained four CPUs. See Appendix A for additional server specifications details.

We tested each server configuration until the CPU utilization across all VMs reached 100% and the requests per second reached a sustained peak level for each server. For 1 and 2 VMs per CPU, we stopped testing at the 60 and 64 WebBench client load respectively. For 3 VMs per CPU, we increased the WebBench client load count to 120.

We conducted a series of WebBench test runs as described in the test methodology section of this report. We performed two runs of each configuration to verify results consistency and reported the maximum results collected for each configuration.

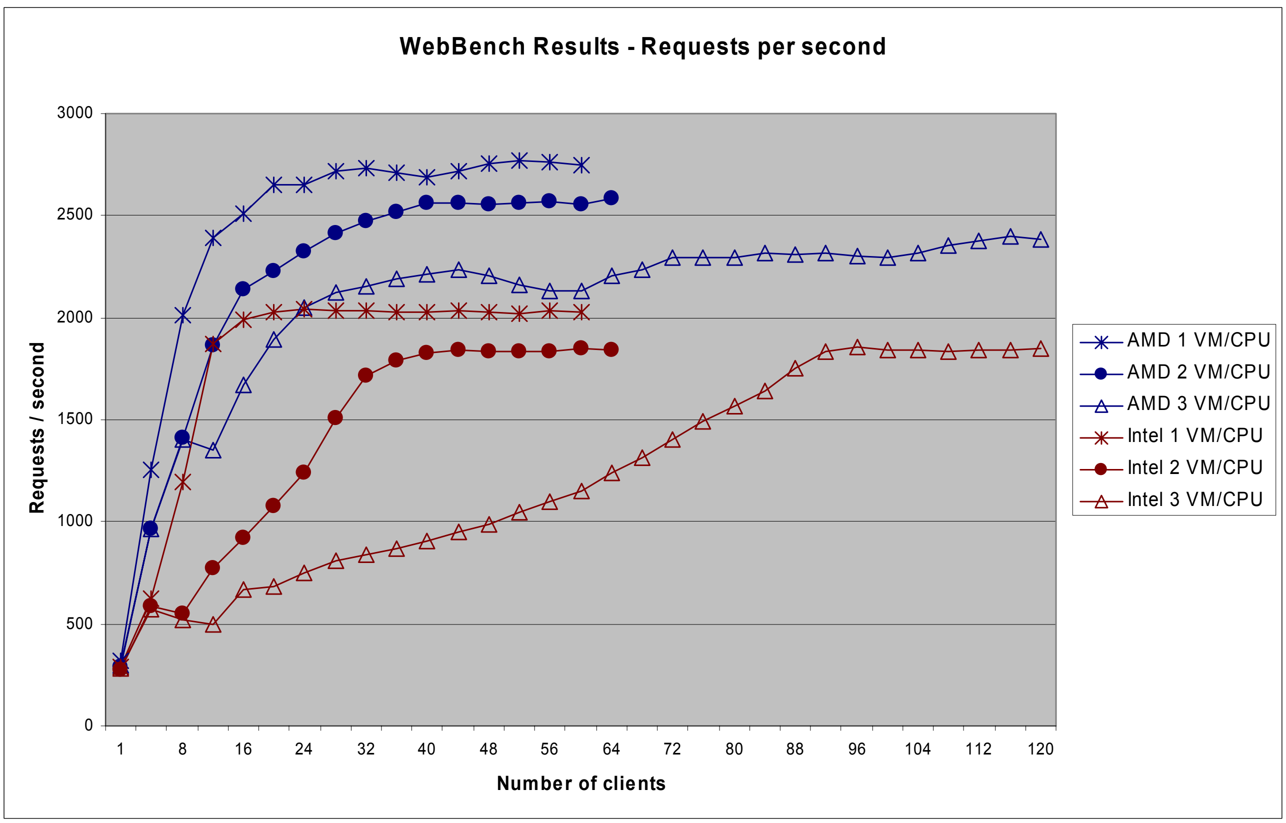

Figure 6: WebBench requests per second results (Greater numbers are better.)

As shown in Figures 6 and 7, we found that the AMD Opteron server running in a 1 VM per CPU configuration had the greatest number of requests per second, with a peak of 2767.79 requests per second in our test environment. The Intel Xeon server running in a 1 VM per CPU configuration had a peak of 2044.87 requests per second, or 26% less than the AMD Opteron 1 VM per CPU configuration.

The AMD Opteron server running in a 2 VM per CPU configuration had a peak of 2580.67 requests per second. The Intel Xeon server running in a 2 VM per CPU configuration had a peak of 1846.85 requests per second, or 28% less than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a peak of 2401.41 requests per second. The Intel Xeon server running in a 3 VM per CPU configuration had a peak of 1854.11 requests per second, or 23% less than the AMD Opteron 3 VM per CPU configuration.

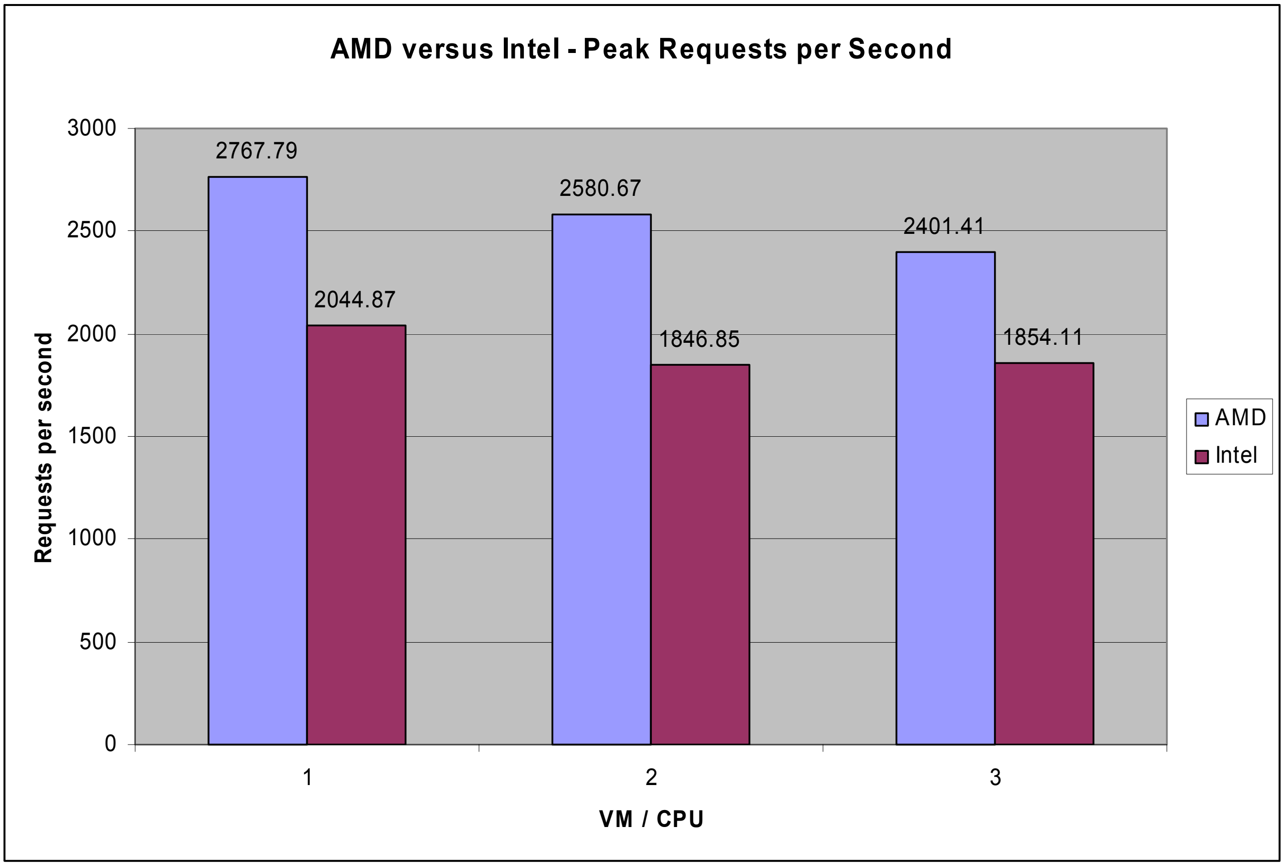

Figure 7: WebBench peak requests per second results (Greater numbers are better.)

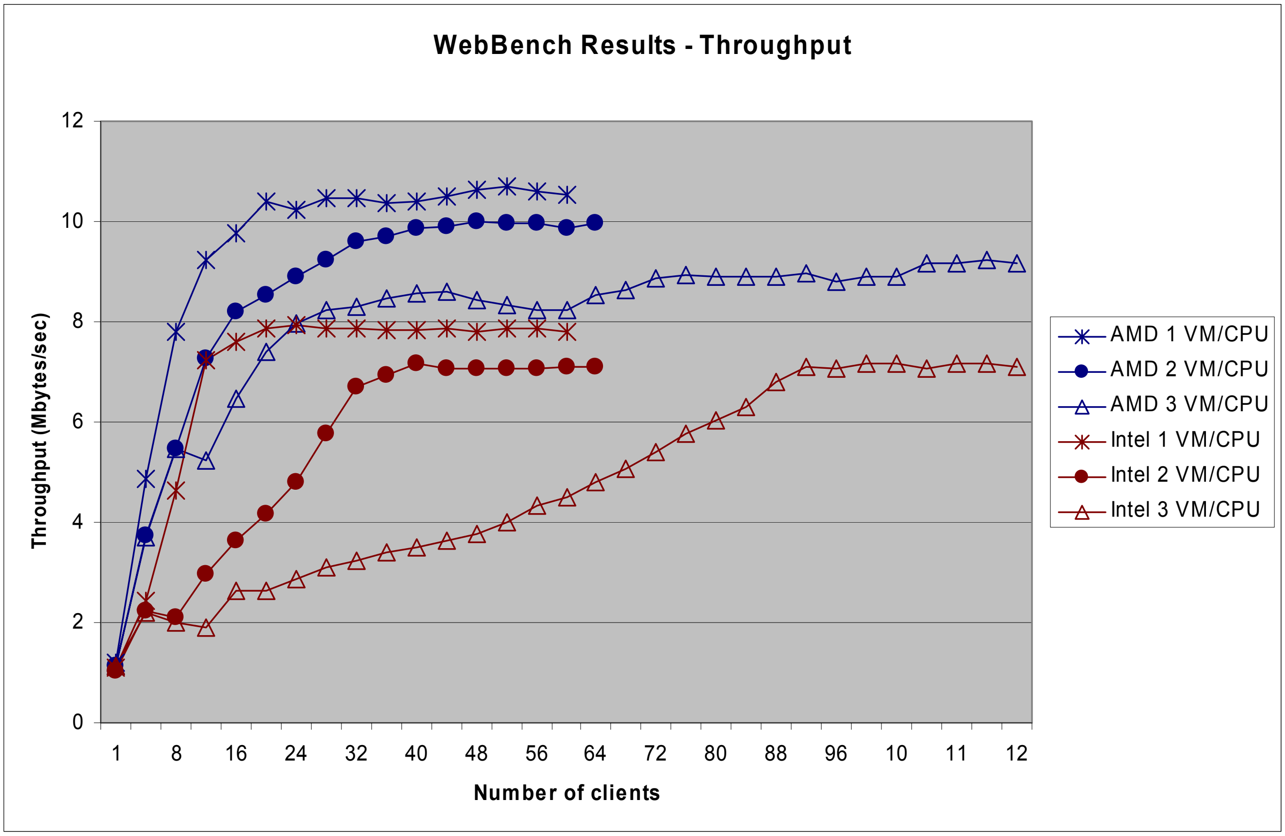

Figure 8: WebBench throughput results (Greater numbers are better.)

As shown in Figures 8 and 9, we found that the AMD Opteron server running in a 1 VM per CPU configuration had the greatest WebBench peak throughput of 10.72 MB per second in our test environment. The Intel Xeon server running in a 1 VM per CPU configuration had a peak throughput of 7.94 MB per second, or 26% less than the AMD Opteron 1 VM per CPU configuration.

The AMD Opteron server running in a 2 VM per CPU configuration had a peak throughput of 10.0 MB per second. The Intel Xeon server running in a 2 VM per CPU configuration had a peak throughput of 7.15 MB per second, or 28% less than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a peak throughput of 9.23 MB per second. The Intel Xeon server running in a 3 VM per CPU configuration had a peak throughput of 7.18 MB per second, or 22% less than the AMD Opteron 3 VM per CPU configuration.

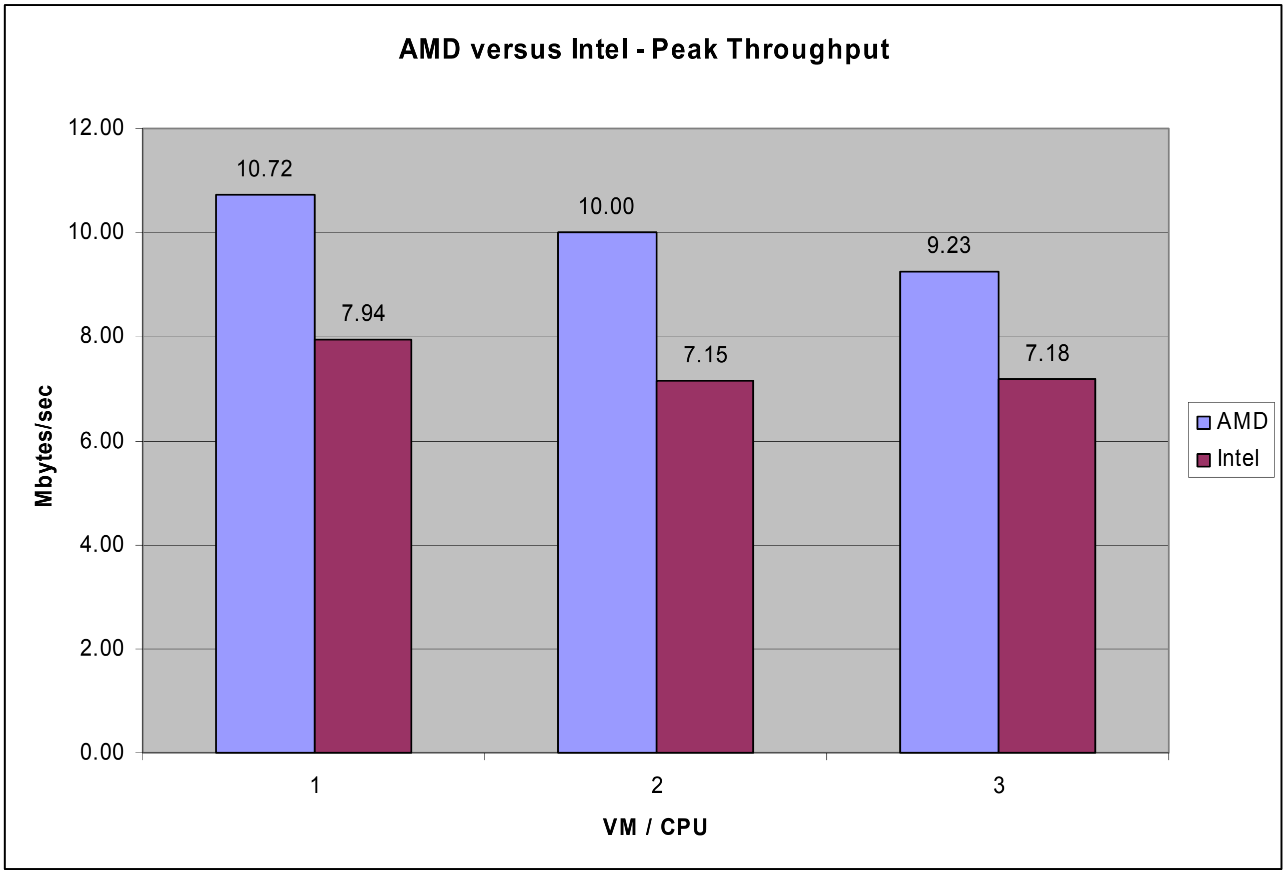

Figure 9: WebBench peak throughput results (Greater numbers are better.)

Overall, we found that the peak requests per second for all VM configurations running on the AMD Opteron server was greater than any of the VM configurations running on Intel Xeon server.

As shown in Figure 10, our results showed that the AMD Opteron server running in 1, 2, and 3 VMs per CPU configurations had lower mean WebBench request latencies than the Intel Xeon server running in 1, 2, and 3 VMs per CPU configurations.

We found that the AMD Opteron server running in a 1 VM per CPU configuration had the lowest WebBench request latency of 65.6 milliseconds at the 60-client load in our test environment. The Intel Xeon server running in a 1 VM per CPU configuration had a request latency of 88.6 milliseconds at the 60-client load, or 35% higher than the AMD Opteron 1 VM per CPU configuration. Lower numbers are better.

The AMD Opteron server running in a 2 VM per CPU configuration had a request latency of 74.5 millisecond at the 64-client load. The Intel Xeon server running in a 2 VM per CPU configuration had a request latency of 110.4 milliseconds at the 64-client load, or 48% higher than the AMD Opteron 2 VM per CPU configuration.

The AMD Opteron server running in a 3 VM per CPU configuration had a request latency of 151.2 milliseconds at the 120-client load. The Intel Xeon server running in a 3 VM per CPU configuration had a request latency of 194.5 milliseconds at the 120-client load, or 29% higher than the AMD Opteron 3 VM per CPU configuration.

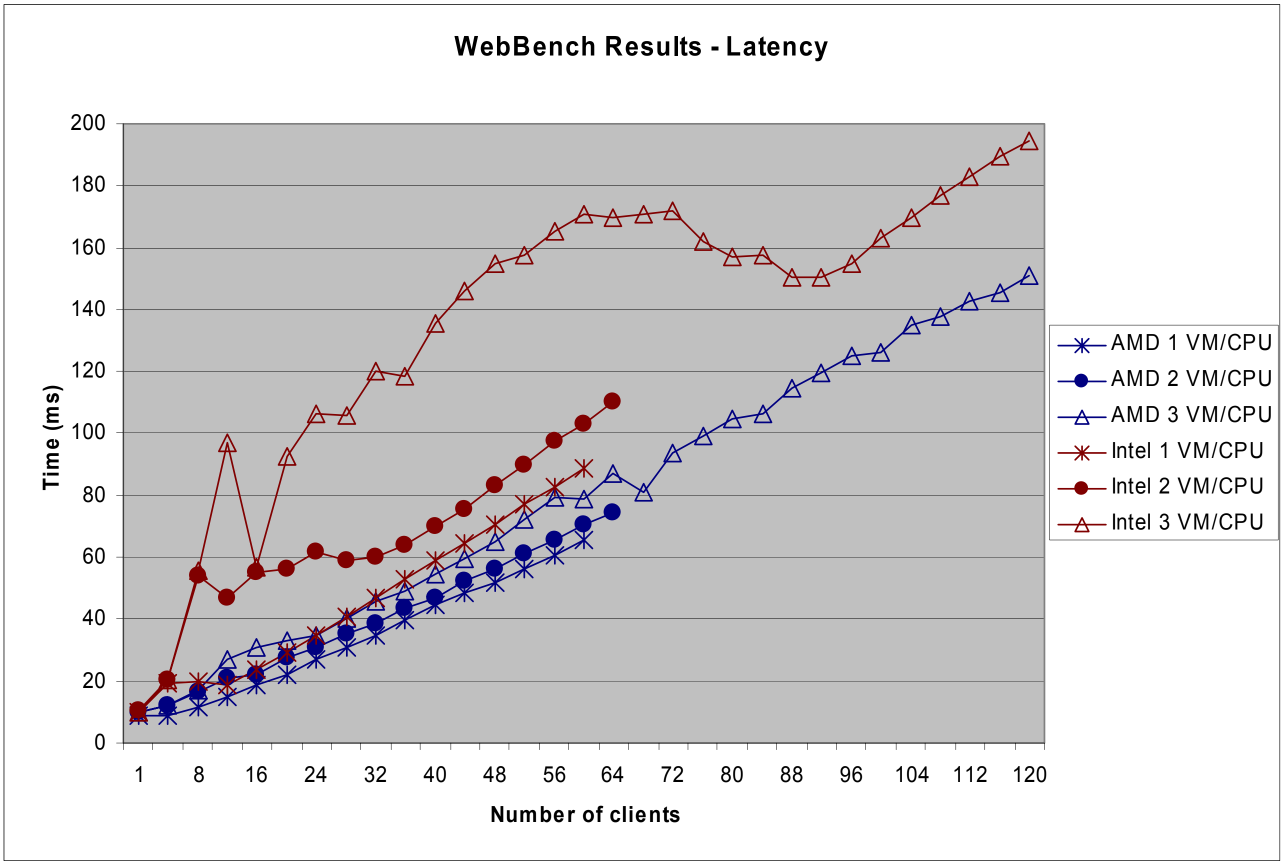

Figure 10: WebBench request latency results (Lower numbers are better.)

None of the WebBench request latencies results that we collected reached an observed maximum value, so we were unable to report peak latencies for this study.

PC2100 versus PC2700

The HP ProLiant HP585 AMD Opteron processor-based server shipped with PC2700 RAM whereas the IBM xSeries 445 Intel Xeon processor-based server shipped with PC2100. During the course of this testing, we replaced the PC2100 RAM in the IBM xSeries 445 server with PC2700 RAM and retested all IBM xSeries 445 test configurations. In our testing, we found less than 2% difference in all collected WebBench results.

Appendices

A. Server and client system disclosure

IBM xSeries 445	
Machine Type	IBM xSeries 445
Processor(s)	4 x 3.00GHz Intel Xeon processors with HyperThreading enabled
Host Adapter	LSI Logic PCI-X Ultra320 SCSI Host Adapter
L1 cache	20 KB
L2 cache	512 KB
L3 cache	4 MB
Hard Drive Controller	ServerRAID-6M
Memory	16GB PC2100 and PC2700
Network Adapter(s)	4 x Intel 82543 Gigabit Ethernet Adapters, Rev 02
Web Server Software	Internet Information Services (IIS) 6.0
Virtualization software	VMware ESX Server 2.5.0 build-11548
Guest OS	Microsoft Windows Server 2003, Enterprise Edition, Version 5.2.3790 Build 3790

Figure A1: IBM xSeries 445 Server Configuration

For details related to the IBM xSeries 445 server, see http://www1.ibm.com/servers/eserver/xseries/x445.html.

HP ProLiant HP585	
Machine Type	HP ProLiant HP585
Processor(s)	4 x 2.40GHz AMD Opteron processors
L1 cache	128 KB
L2 cache	1 MB
L3 cache	NA
Hard Drive Controller	HP Smart Array 5i disk controller
Memory	16GB PC2700
Network Adapter(s)	4 x Intel 82543 Gigabit Ethernet Adapters, Rev 02
Web Server Software	Internet Information Services (IIS) 6.0
Virtualization software	VMware ESX Server 2.5.0 build-11548
Guest OS	Microsoft Windows Server 2003, Enterprise Edition, Version 5.2.3790 Build 3790

Figure A2: HP ProLiant HP585 Server Configuration

For details related to the HP ProLiant DL585 server, see http://h18004.www1.hp.com/products/servers/proliantdl585/index.html.

Network Test Bed Clients	
Machine Type	Dell PowerEdge 350
BIOS	Dell TR440BXA.86B.0034.P12
Processor(s)	1 – 850MHz Pentium III
Hard Drive	Maxtor 5T010H1 10GB
Memory	256MB
L2 Cache	256K
Network Adapter(s)	Intel Pro100+ Management Adapter
Video Card	ATI Rage
OS	Windows XP/SP1 / Windows 2000 Professional/SP4

Figure A3: Network Test Bed Clients Configuration

B. Testbed configuration

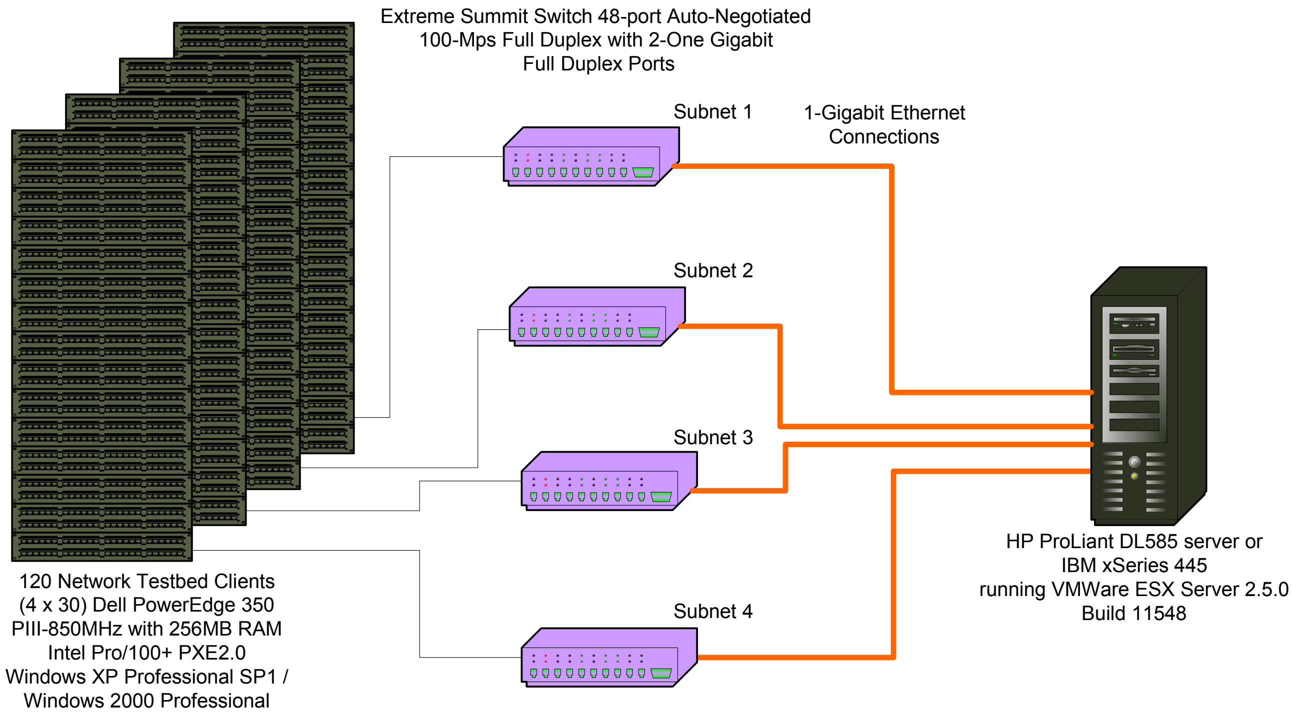

Figure B1: WebBench Network Configuration

We designed the testbed environment to distribute the WebBench workload across all of the subnets in an evenly distributed manner.

For the **60-client test**, we used IP addresses 200.105.1.1 through 200.105.1.15 on subnet one connected to NIC one. We used IP addresses 200.105.2.31 through 200.105.2.45 on subnet two connected to NIC two. We used IP addresses 200.106.1.1 through 200.106.1.15 on subnet three connected to NIC three. We used IP addresses 200.106.2.31 through 200.106.2.45 on subnet four connected to NIC four.

For the **64-client test**, we used IP addresses 200.105.1.1 through 200.105.1.16 on subnet one connected to NIC one. We used IP addresses 200.105.2.31 through 200.105.2.46 on subnet two connected to NIC two. We used IP addresses 200.106.1.1 through 200.106.1.16 on subnet three connected to NIC three. We used IP addresses 200.106.2.31 through 200.106.2.46 on subnet four connected to NIC four.

For the **120-client test**, we used IP addresses 200.105.1.1 through 200.105.1.30 on subnet one connected to NIC one. We used IP addresses 200.105.2.31 through 200.105.2.60 on subnet two connected to NIC two. We used IP addresses 200.106.1.1 through 200.106.1.30 on subnet three connected to NIC three. We used IP addresses 200.106.2.31 through 200.106.2.60 on subnet four connected to NIC four.

C. WebBench client to VM mapping

WebBench client allocation in the 4 VM.

1 Clients

Clients	VM
1	1
	2
	3
	4

For the 1-client workload, only VM 1 receives load.

4 Clients

Clients	VM
1	1
2	2
3	3
4	4

For the 4-client workload, each VM receives 1 client load.

8 Clients

Clients	Clients	VM
1	5	1
2	6	2
3	7	3
4	8	4

12 Clients

Clients	Clients	Clients	VM
1	5	9	1
2	6	10	2
3	7	11	3
4	8	12	4

16 clients

Clients	Clients	Clients	Clients	VM
1	5	9	13	1
2	6	10	14	2
3	7	11	15	3
4	8	12	16	**4**

20 clients

Clients	Clients	Clients	Clients	Clients	VM
1	5	9	13	17	**1**
2	6	10	14	18	**2**
3	7	11	15	19	**3**
4	8	12	16	20	**4**

32 clients

C	C	C	C	C	C	C	C	VM
1	5	9	13	17	21	25	29	**1**
2	6	10	14	18	22	26	30	**2**
3	7	11	15	19	23	27	31	**3**
4	8	12	16	20	24	28	32	**4**

64 clients

C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	VM
1	5	9	13	17	21	25	29	33	37	41	45	49	53	57	62	**1**
2	6	10	14	18	22	26	30	34	38	42	46	50	54	58	62	**2**
3	7	11	15	19	23	27	31	35	39	43	47	51	55	59	63	**3**
4	8	12	16	20	24	28	32	36	40	44	48	52	56	60	64	**4**

When the server is configured with four VMs, then each VM receives 8 WebBench clients at the 64 concurrent client level.

WebBench client allocation in the 12 VM.

1 Client

Clients	VMs
1	**1**
	2
	3
	4
	5
	6
	7
	8
	9
	10
	11
	12

For the 1-client workload, only VM 1 receives load.

16 Clients

Clients	Clients	VMs
1	13	**1**
2	14	**2**
3	15	**3**
4	16	**4**
5		**5**
6		**6**
7		**7**
8		**8**
9		**9**
10		**10**
11		**11**
12		**12**

For the 16-client workload, most VMs receive a 1-client load and the first 4 VMs receive a 2-client load.

32 Clients

Clients	Clients	Clients	VMs
1	13	25	**1**
2	14	26	**2**
3	15	27	**3**
4	16	28	**4**
5	17	29	**5**
6	18	30	**6**
7	19	31	**7**
8	20	32	**8**
9	21		**9**
10	22		**10**
11	23		**11**
12	24		**12**

48 Clients

Clients	Clients	Clients	Clients	VMs
1	13	25	37	**1**
2	14	26	38	**2**
3	15	27	39	**3**
4	16	28	40	**4**
5	17	29	41	**5**
6	18	30	42	**6**
7	19	31	43	**7**
8	20	32	44	**8**
9	21	33	45	**9**
10	22	34	46	**10**
11	23	35	47	**11**
12	24	36	48	**12**

120 Clients

Clients	Clients	Clients	Clients	Clients	Clients	Clients	Clients	Clients	Clients	VMs
1	13	25	37	49	61	73	85	97	109	**1**
2	14	26	38	50	62	74	86	98	110	**2**
3	15	27	39	51	63	75	87	99	111	**3**
4	16	28	40	52	64	76	88	100	112	**4**
5	17	29	41	53	65	77	89	101	113	**5**
6	18	30	42	54	66	78	90	102	114	**6**
7	19	31	43	55	67	79	91	103	115	**7**
8	20	32	44	56	68	80	92	104	116	**8**
9	21	33	45	57	69	81	93	105	117	**9**
10	22	34	46	58	70	82	94	106	118	**10**
11	23	35	47	59	71	83	95	107	119	**11**
12	24	36	48	60	72	84	96	108	120	**12**

When the server is configured with 12 VMs, then each VM receives 10 WebBench clients at the 120 concurrent client level.

VeriTest (www.veritest.com), the testing division of Lionbridge Technologies, Inc., provides outsourced testing solutions that maximize revenue and reduce costs for our clients. For companies who use high-tech products as well as those who produce them, smoothly functioning technology is essential to business success. VeriTest helps our clients identify and correct technology problems in their products and in their line of business applications by providing the widest range of testing services available.

VeriTest created the suite of industry-standard benchmark software that includes WebBench, NetBench, Winstone, and WinBench. We've distributed over 20 million copies of these tools, which are in use at every one of the 2001 Fortune 100 companies. Our Internet BenchMark service provides the definitive ratings for Internet Service Providers in the US, Canada, and the UK.

Under our former names of ZD Labs and eTesting Labs, and as part of VeriTest since July of 2002, we have delivered rigorous, objective, independent testing and analysis for over a decade. With the most knowledgeable staff in the business, testing facilities around the world, and almost 1,600 dedicated network PCs, VeriTest offers our clients the expertise and equipment necessary to meet all their testing needs.

For more information email us at info@veritest.com or call us at 919-380-2800.